arXiv:1611.06906v1 [cs.CV] 21 Nov 2016

Multi-Scale Anisotropic Fourth-Order Diffusion Improves Ridge and Valley Localization

Shekoufeh Gorgi Zadeh[*], Stephan Didas[†], Maximilian W. M. Wintergerst[‡], and Thomas Schultz[*]

November 22, 2016

Ridge and valley enhancing filters are widely used in applications such as vessel detection in medical image computing. When images are degraded by noise or include vessels at different scales, such filters are an essential step for meaningful and stable vessel localization. In this work, we propose a novel multi-scale anisotropic fourth-order diffusion equation that allows us to smooth along vessels, while sharpening them in the orthogonal direction. The proposed filter uses a fourth order diffusion tensor whose eigentensors and eigenvalues are determined from the local Hessian matrix, at a scale that is automatically selected for each pixel. We discuss efficient implementation using a Fast Explicit Diffusion scheme and demonstrate results on synthetic images and vessels in fundus images. Compared to previous isotropic and anisotropic fourth-order filters, as well as established second-order vessel enhancing filters, our newly proposed one better restores the centerlines in all cases.

1 Introduction

In image analysis, ridges and valleys are curves along which the image is brighter or darker, respectively, than the local background [7]. Collectively, ridges and valleys are referred to as creases. Reliable detection and localization of creases in noisy images is an important and well-studied problem in medical image analysis, one very common application being the detection of blood vessels [8].

[*]Institute of Computer Science II, University of Bonn, Germany
[†]Umwelt-Campus Birkenfeld, Hochschule Trier, Germany
[‡]Department of Ophthalmology, University of Bonn, Germany

Figure 1: On smoothly shaded surfaces, second-order Perona-Malik diffusion creates a staircasing artifact that is avoided by fourth-order diffusion.

Often, ridges and valleys occur at multiple scales, for example due to the branching of vessel trees. Since Gaussian scale space [16] does not offer any specific mechanisms for preserving creases, image filters such as coherence enhancing diffusion [22], crease enhancement diffusion (CED) [21], and vesselness enhancement diffusion (VED) [4] have become popular. They are based on second order anisotropic diffusion equations with a diffusion tensor that, in the presence of crease lines, smoothes only along, but not across them. In addition, the VED filter includes a multi-scale analysis that automatically adapts it to the local scale of creases.

In this work, we argue that using fourth-order instead of second-order diffusion to enhance creases allows for a more accurate localization of their centerlines. We propose a novel fourth-order filter that introduces a fourth-order diffusion tensor to specifically enhance ridges, valleys, or both, in a scale-adaptive manner. Increased accuracy of the final segmentation is demonstrated on simulated and real-world medical images.

2 Related Work

Despite the long history of research in this area, improved filtering and detection of ridges continues to be an active topic in medical image analysis. Our work on improving localization through fourth-order diffusion complements recent advances. For example, the SCIRD ridge detector by Annunziata et al. [1], or the vesselness measure by Jerman et al. [13] that gives better responses for vessels of varying contrasts, could replace the vessel segmentation by Frangi et al. [8] that we use as a prefiltering step, or the work by Hannink et al. [10] could be used to improve the performance of our filter in crossings and bifurcations.

Fourth-order diffusion generalizes the heat equation by replacing all first-order spatial derivatives with second-order derivatives. Our new filter is based on a nonlinear fourth-order diffusion equation that was introduced by Lysaker et al. [17] with the primary goal of avoiding the staircasing artifacts that occur in second-order Perona-Malik diffusion

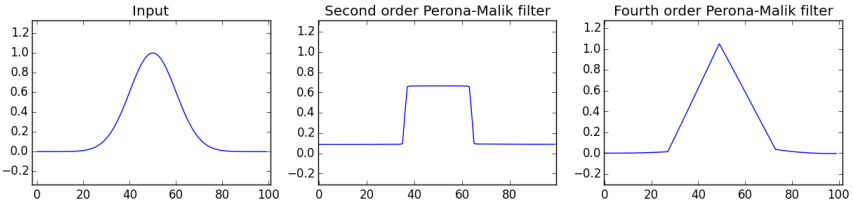

Figure 2: Second-order diffusion filter and fourth-order diffusion filter applied on a 1-D input signal. Both filters use the Perona-Malik diffusivity function.

[18], as shown on the "Pepper" image in Figure 1. Subsequently, Didas et al. [6] have shown that, when combined with specific diffusivity functions, such fourth-order equations can enhance image curvature analogous to how, by careful use of forward and backward diffusion, second-order equations can enhance, rather than just preserve, image edges [22].

Figure 2 shows a simple one-dimensional example with the Perona-Malik diffusivity function. The edge enhancement of nonlinear second-order diffusion, which leads to a piecewise constant result, and the curvature enhancement of nonlinear fourth-order diffusion, which leads to a piecewise linear result, are clearly visible. Obviously, localizing the maximum will be much easier and more reliable in case of the sharp peak created by fourth-order diffusion than in the extended plateau that results from second-order diffusion.

It is this curvature-enhancing property of fourth-order diffusion that we exploit in our novel filter. In particular, we introduce a novel anisotropic fourth-order diffusion equation that smoothes along the crease, while creating a sharp peak in the orthogonal direction, to clearly indicate its center. We are aware of only one previous formulation of anisotropic fourth order diffusion, proposed by Hajiaboli [9]. However, it has been designed to preserve edges, rather than enhance creases. Consequently, it is not well-suited for our purposes, as we will demonstrate in the results. Moreover, it differs from our approach in that it does not make use of a fourth-order diffusion tensor, and includes no mechanism for scale selection.

3 Method

3.1 Anisotropic Fourth-order Diffusion

Diffusion-based image filters treat image intensities as an initial heat distribution $u_{t=0}$, and solve the heat equation $\partial_t u = \operatorname{div}(g \nabla_{\mathbf{x}} u)$ for larger values of an artificial time parameter t, corresponding to increasingly smoothed versions of the image. When the scalar diffusivity g is made a function of the spatial gradient magnitude $\|\nabla_{\mathbf{x}} u\|$, diffusion becomes nonlinear, and image edges can be preserved by reducing the amount of smoothing near them [18].

Anisotropic diffusion $\partial_t u = \operatorname{div}(D \nabla_{\mathbf{x}} u)$ replaces the scalar diffusivity g by a second-order diffusion tensor D, which makes it possible to reduce smoothing orthogonal to, but

not along image features, and therefore to denoise edges more effectively than isotropic nonlinear diffusion, while still avoiding to destroy them [22]. *Fourth-order* filters generalize the diffusion equation in a different way, by replacing the two first-order spatial derivatives in the heat equation with second-order derivatives, leading to an overall order of four. In particular, building on work of Lysaker et al. [17], Didas et al. [6] formulate nonlinear fourth-order diffusion as

$$\begin{aligned}\partial_t u = &- \partial_{xx}(g(\|H(u)\|_F^2)u_{xx}) - \partial_{yx}(g(\|H(u)\|_F^2)u_{xy}) \\ &- \partial_{xy}(g(\|H(u)\|_F^2)u_{yx}) - \partial_{yy}(g(\|H(u)\|_F^2)u_{yy}).\end{aligned} \tag{1}$$

We propose the following novel *anisotropic fourth-order* diffusion model, which combines the ideas of higher-order diffusion with that of making diffusivity a function of both spatial location and direction:

$$\begin{aligned}\partial_t u = &- \partial_{xx}\left[\mathcal{D}(H_\rho(u_\sigma)) : H(u)\right]_{xx} - \partial_{yx}\left[\mathcal{D}(H_\rho(u_\sigma)) : H(u)\right]_{xy} \\ &- \partial_{xy}\left[\mathcal{D}(H_\rho(u_\sigma)) : H(u)\right]_{yx} - \partial_{yy}\left[\mathcal{D}(H_\rho(u_\sigma)) : H(u)\right]_{yy}\end{aligned} \tag{2}$$

Equation (2) introduces a general linear map $\mathcal{D}$ from the Hessian matrix H to a transformed matrix. Linear maps from matrices to matrices are naturally written as fourth-order tensors, and we use the "double dot product" $\mathcal{D} : H$ as a shorthand for applying the map $\mathcal{D}$ to the Hessian matrix H. This results in a transformed matrix T, and we use square brackets $[T]_{ij}$ to denote its (i, j)th component. Formally,

$$[T]_{ij} = \left[\mathcal{D}(H_\rho(u_\sigma)) : H(u)\right]_{ij} = \sum_{k=1}^{2}\sum_{l=1}^{2}\left[\mathcal{D}(H_\rho(u_\sigma))\right]_{ijkl}\left[H(u)\right]_{kl}. \tag{3}$$

In this notation, we can define second-order eigentensors E of $\mathcal{D}$ corresponding to eigenvalue μ by the equation $\mathcal{D} : E = \mu E$. An alternative notation, which will be used for the numerical implementation in Section 3.4, writes the Hessian and transformed matrices as vectors. This turns $\mathcal{D}$ into a matrix whose eigenvectors are nothing but the vectorized eigentensors as defined above. Similar to others [3, 14], we find the fourth-order tensor and "double dot" notation more appealing for reasoning at a higher level, because it allows us to preserve the natural structure of the involved matrices.

Using our square bracket notation, an equivalent way of writing one of the terms from Equation (1), $\partial_{ji}(g(\|H(u)\|_F^2)u_{ij})$, is $\partial_{ji}\left[g(\|H(u)\|_F^2)H(u)\right]_{ij}$. This clarifies that the difference between the model from Equation (1) and our new one in Equation (2) is to replace the isotropic scaling of Hessian matrices that uses a scalar diffusivity g, with a general linear transformation $\mathcal{D}$, which acts on the second-order Hessian in analogy to how the established second-order diffusion tensor acts on gradients in second-order anisotropic diffusion. Due to this analogy, we call $\mathcal{D}$ a fourth-order diffusion tensor.

In our filter, $\mathcal{D}$ is a function of the local normalized Hessians, which are defined as

$$H_\rho(u_\sigma) = G_\rho * \left(\frac{1}{\sqrt{1 + \|\nabla u_\sigma\|}} H\left(u_\sigma\right)\right), \tag{4}$$

where regularized derivatives are obtained by convolution with a Gaussian kernel, $u_\sigma := u * G_\sigma$. Its width σ should reflect the scale of the crease, as will be discussed in Section 3.3. Since scale selection might introduce spatial discontinuities in the chosen σ, the normalized Hessians are made differentiable by integrating them over a neighborhood, for which we use a Gaussian width $\rho = 0.5$ in our experiments. As shown in [11], and used for vesselness enhancement diffusion in [4], the inverse gradient magnitude factor is used to make the eigenvalues of $H_\rho(u_\sigma)$ match the surface curvature values.

We emphasize that, unlike in a previous generalization of structure tensors to higher order [20], the reason for going to higher tensor order in Equation (2) is not to preserve information at crossings; this is a separate issue that was recently addressed by others [10], and that we plan to tackle in our own future work. In our present work, our goal is to smooth along ridges and valleys, while sharpening them in the orthogonal direction. This sharpening requires the curvature-enhancing properties of fourth-order diffusion, and a fourth-order diffusion tensor is a natural consequence of making it anisotropic.

3.2 Fourth-order Diffusion Tensor $\mathcal{D}$

We now need to construct our fourth-order diffusion tensor $\mathcal{D}$ so that it will smooth along creases, while enhancing them in the perpendicular direction. Similar to Weickert's diffusion tensors [22], we will construct $\mathcal{D}$ in terms of its eigentensors E_i and corresponding eigenvalues μ_i, as defined above.

Didas et al. [6] have shown that fourth-order diffusion with the Perona-Malik diffusivity [18] allows for adaptive smoothing or sharpening of image curvature, depending on a contrast parameter λ. In particular, in the 1-D case, only forward diffusion (i.e., smoothing) happens in regions with $|\partial_{xx}u| < \lambda$, while only backward diffusion (i.e., curvature enhancement) occurs where $|\partial_{xx}u| > \sqrt{3}\lambda$. We wish to exploit this to enhance creases whose curvature is strong enough to begin with, while smoothing out less significant image features.

This is achieved by deriving the eigenvalues μ_i of $\mathcal{D}$ from the eigenvalues ν_1, ν_2 of the normalized Hessian $H_\rho(u_\sigma)$ using the Perona-Malik diffusivity [18], i.e., $\mu_i = (1 + \nu_i^2/\lambda^2)^{-1}$ for $i \in \{1, 2\}$. If the user wishes to specifically enhance either ridges or valleys, the sign of ν_i must be taken into account. For instance, a ridge-like behaviour in the ith direction is characterized by $\nu_i < 0$. Therefore, we can decide to smooth out valleys by setting $\mu_i = 1$ wherever $\nu_i \geq 0$, and enhance ridges wherever $\nu_i < 0$ by defining μ_i as before. Enhancing only valleys can be done in full analogy.

The ridge and valley directions can be found from the eigenvectors e_1, e_2 of the normalized Hessian matrix $H_\rho(u_\sigma)$, and are reflected in the eigentensors E_i of $\mathcal{D}$ by setting

$$\begin{array}{lll} E_1 = & e_1 \otimes e_1 & \qquad E_3 = \frac{1}{\sqrt{2}}(e_1 \otimes e_2 + e_2 \otimes e_1) \\ E_2 = & e_2 \otimes e_2 & \qquad E_4 = \frac{1}{\sqrt{2}}(e_1 \otimes e_2 - e_2 \otimes e_1) \end{array} \quad . \tag{5}$$

The E_i are orthonormal with respect to the tensor dot product $A : B := \mathrm{tr}(B^{\mathrm{T}} A)$. By definition, E_4 is antisymmetric. Since Hessians of smooth functions are symmetric, the value of μ_4 does not play a role, and is simply set to zero. We define μ_3 as the average of μ_1 and μ_2.

3.3 Scale Selection

Similar to vesselness enhancement diffusion (VED) [4], our approach uses the vesselness measure by Frangi et al. [8] to automatically select the most suitable local scale σ. It is obtained from sorted and scale-normalized eigenvalues $|\tilde{\nu}_1| \leq |\tilde{\nu}_2|$, computed as $\tilde{\nu}_i := \sigma^2 \nu_i$ from eigenvalues ν_i of the Hessian $H(u_\sigma)$ at a given scale σ. The factor σ^2 compensates for the loss of contrast at larger scales [16].

A vesselness measure $\mathcal{V}_\sigma$ should be low in background regions where overall curvature and thus $\mathcal{S} = \sqrt{\tilde{\nu}_1^2 + \tilde{\nu}_2^2}$ are low overall. Moreover, it should detect tubular structures, where $|\tilde{\nu}_1| \ll |\tilde{\nu}_2|$, as opposed to blobs, in which $\mathcal{R}_\mathcal{B} = \frac{\tilde{\nu}_1}{\tilde{\nu}_2}$ would be large. For ridges ($\tilde{\nu}_2 < 0$), this is achieved by combining $\mathcal{S}$ and $\mathcal{R}_\mathcal{B}$ according to

$$\mathcal{V}_\sigma u = \begin{cases} 0 & \text{if} \quad \tilde{\nu}_2 > 0 \\ \left(\mathrm{e}^{-\frac{\mathcal{R}_\mathcal{B}^2}{2\beta^2}}\right)\left(1 - \mathrm{e}^{-\frac{\mathcal{S}^2}{2c^2}}\right) & \text{otherwise} \end{cases}, \tag{6}$$

where the β and c parameters tune $\mathcal{V}_\sigma u$ to be more specific with respect to suppression of blob shapes or background structures, respectively. We use $\beta = 0.5$ and $c = \frac{1}{2}(\max(\mathcal{S}))$, as recommended by Frangi et al. [8].

At each pixel, vesselness should be maximal at the scale that matches the corresponding vessel size. Therefore, the scale for each pixel can be selected as the σ for which the maximum $\mathcal{V}u = \max_{\sigma=\sigma_{\min},\ldots,\sigma_{\max}} \mathcal{V}_\sigma u$ is attained, where $\{\sigma_{\min}, \ldots, \sigma_{\max}\}$ are the range of expected scales in the image. For pixels that are part of the background, $\mathcal{V}u$ is low, and it can be thresholded by parameter $\theta \in [0, 1]$ for vessel segmentation.

It has been observed previously [5] that vesselness often fails to correctly estimate the scale of the vessel along its boundary. This can happen in two cases: When the ridge has a step-like shape, the curvature near the corner points will be much larger at the finest scale than at all other scales, leading to an underestimation of the real scale near the boundary. On the other hand, the cross-sectional intensity profile of vessels may have inflection points near its edges, where $\tilde{\nu}_2$ changes its sign. In this case, some points near the boundary will have zero vesselness at the finest scale, but the sign of $\tilde{\nu}_2$ will flip, and therefore vesselness becomes non-zero, at coarser scales, leading to an overestimation of scale. Figure. 3 shows both scale underestimation or overestimation at vessel boundaries.

Such boundary effects are less problematic for VED, but can lead to serious artifacts when vesselness is used to steer curvature-enhancing diffusion. In particular, they can cause our filter to enhance the boundary of large-scale ridges more than their center. We post process the computed scales to avoid this. For each pixel on a vessel, the vessel cross-section containing that pixel is extracted by following the eigenvector direction that corresponds to the strongest eigenvalue of the Hessian matrix computed at the scale suggested by the vesselness measures at each point. After post-processing, all pixels are assigned the scale closest to the average of all pixels that lie on the same cross-section. This removes the problem of scale overestimation or underestimation on the boundaries. Figure. 3 shows the scale image before and after being post processed.

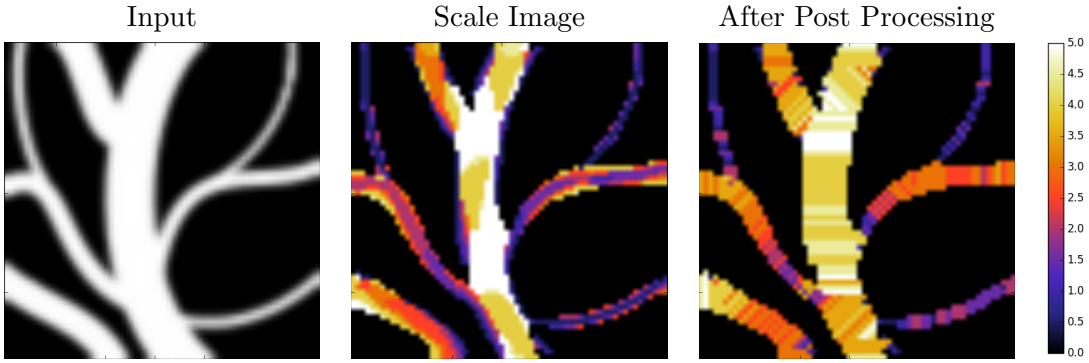

Figure 3: The color of each pixel in the scale image illustrates the scale of the underlying vessel. Artifacts near the boundaries are removed in a post processing step.

3.4 Stability

In order to solve Equation (2), we discretize it with standard finite differences, and use an explicit numerical scheme. In matrix-vector notation, this can be written as

$$u^{k+1} = u^k - \tau P u^k = (I - \tau P) u^k \ , \tag{7}$$

where $u^k \in \mathbb{R}^m$ is the vectorized image at iteration k, and the exact form of matrix $P \in \mathbb{R}^{m \times m}$ will be discussed later. We call a numerical scheme ℓ_2 stable if

$$\left\| u^{k+1} \right\|_2 \leq \left\| u^k \right\|_2 \ , \tag{8}$$

i.e., the ℓ_2 norm of the image is guaranteed not to increase from iteration k to $k+1$. It follows from Equation (7) that

$$\left\| u^{k+1} \right\|_2 \leq \left\| I - \tau P \right\|_2 \cdot \left\| u^k \right\|_2 \ , \tag{9}$$

where $\|P\|_2$ denotes the ℓ_2 norm of P, i.e., $\|P\|_2 := \sqrt{\rho(P^T P)}$, where $\rho(P^T P)$ computes the largest modulus of eigenvalues of the symmetric matrix $P^T P$.

Consequently, the condition in Equation (9) is satisfied if

$$\left\| I - \tau P \right\|_2 \leq 1 \ . \tag{10}$$

Since P is positive semi-definite, the eigenvalues of $I - \tau P$ are within the interval $[1 - \tau \left\| P \right\|_2, 1]$. Thus, Equation (10) is satisfied if $1 - \tau \left\| P \right\|_2 \geq -1$. This results in the following constraint on the permissible time step size τ:

$$\tau \leq \frac{2}{\left\| P \right\|_2} \tag{11}$$

This clarifies that the restriction on the time step size only depends on $\|P\|_2$. To compute it, we will now write down the system matrix P for our discretization of fourth-order anisotropic diffusion filtering.

Let L_{xx}, L_{xy}, L_{yx}, L_{yy} be the matrices approximating the corresponding derivatives. For "natural" boundary condition, it is important only to approximate the derivatives at pixels where the whole stencil fits in the image domain, i.e., where enough data is available. Let us combine these four matrices pixelwise into one big matrix L such that

$$Lu \approx \begin{pmatrix} \vdots \\ L_{xx_i} u_i \\ L_{xy_i} u_i \\ L_{yx_i} u_i \\ L_{yy_i} u_i \\ \vdots \end{pmatrix}, \tag{12}$$

i.e., the approximations of the four derivatives will be next to each other for every pixel.

The 4×4 matrix form of the fourth-order diffusion tensor in pixel i, acting on $\left(L_{xx_i} u_i \; L_{xy_i} u_i \; L_{yx_i} u_i \; L_{yy_i} u_i\right)^T$, can be written as $D_i = EME^T$, where E is an orthogonal matrix containing the vectorized E_1, E_2, E_3, E_4 from Equation (5) as its columns and M is a diagonal matrix with the eigenvalues μ_1, μ_2, μ_3, μ_4 on its diagonal. Due to the choice of the Perona-Malik diffusivity in our model, $\|D_i\|_2 \leq 1$.

If we arrange all per-pixel matrices D_i in one big matrix D with a 4×4 block-diagonal structure,

$$D = \begin{pmatrix} [D_1] & & \cdots & & 0 \\ & [D_2] & & & \\ \vdots & & [D_3] & & \vdots \\ & & & \ddots & \\ 0 & & \cdots & & [D_m] \end{pmatrix}, \tag{13}$$

it is clear that $\|D\|_2 \leq 1$, and the whole scheme reads as

$$u^{k+1} = u^k - \tau L^T D L u^k. \tag{14}$$

Substituting into Equation (11) yields

$$\tau \leq \frac{2}{\|L^T D L\|_2}, \tag{15}$$

meaning that, in order to find a stable step size τ, we have to bound

$$\left\|L^T D L\right\|_2 \leq \|L\|_2^2 \leq \|L_{xx}\|_2^2 + \|L_{xy}\|_2^2 + \|L_{yx}\|_2^2 + \|L_{yy}\|_2^2, \tag{16}$$

whose value will depend on the exact second-order finite difference stencils. We will use the same discretization as Hajiaboli [9], i.e.,

$$\begin{aligned} u_{xx} &\approx \left(u_{i-1,j} - 2u_{i,j} + u_{i+1,j}\right) / \left(\Delta x\right)^2 \\ u_{yy} &\approx \left(u_{i,j-1} - 2u_{i,j} + u_{i,j+1}\right) / \left(\Delta y\right)^2 \\ u_{xy} = u_{yx} &\approx \left(u_{i-1,j-1} + u_{i+1,j+1} - u_{i-1,j+1} - u_{i+1,j-1}\right) / \left(4\Delta x \Delta y\right) \; , \end{aligned}$$

where Δx and Δy are the pixel edge lengths in x and y directions, respectively. It is easy to verify using Gershgorin's theorem that this results in

$$\tau \leq \frac{2}{16\,(\Delta x)^2 + 16\,(\Delta y)^2 + 2\,(\Delta x \Delta y)}, \tag{17}$$

i.e., for $\Delta x = \Delta y = 1$, $\tau \leq 1/17$. In our numerical experiments, we set $\tau = 0.05$.

3.5 Implementation Using Fast Explicit Diffusion

Since the time step size τ derived in the previous section is rather small, solving the discretized version of Equation (2) numerically using a simple explicit Euler scheme requires significant computational effort. The recently proposed Fast Explicit Diffusion (FED) provides a considerable speedup by varying time steps in cycles, in a way that up to half the time steps within a cycle can violate the stability criterion, but the cycle as a whole still remains stable [23]. Consequently, a much smaller number of iterations is required to reach the desired stopping time.

The FED scheme is defined as follows:

$$\begin{cases} u^{k+1,0} = u^k \ , \\ u^{k+1,\,i+1} = \left(I - \tau_i P(u_\sigma^k)\right) u^{k+1,\,i} & (i = 0, \ldots, n-1)\,, \\ u^{k+1} = u^{k+1,\,n} \end{cases} \tag{18}$$

where index k is the cycle iterator, i is the inner cycle iterator, and n is the number of sub-steps in each cycle. In order to ensure stability, $P(u_\sigma^k)$ must be constant during each cycle, and the inner iterations must be executed in a specific order.

The fast explicit diffusion framework can be combined with our discretization in a straightforward manner, and has led to a speedup of around two orders of magnitude in some of our real-world examples. In our experiments on synthetic data, we still used the standard Euler scheme with fixed step sizes, as it allows us to find the exact stopping time that minimizes the ℓ_2 difference between the filtered image and the noise-free ground truth.

3.6 Relationship to Ridge Detection with Scale Selection

In Lindeberg's seminal work on detecting ridges in Gaussian scale space [16], automated scale selection is done using a ridge strength measure such as

$$R(u_\sigma) = t^{4\gamma}\,(u_{xx} + u_{yy})^2 \left((u_{xx} - u_{yy})^2 + 4u_{xy}^2\right) \ , \tag{19}$$

where diffusion time t relates to Gaussian width σ ($t = \sigma^2/2$), and the exponent γ in the normalization factor that is used to compensate for the loss of contrast at later diffusion times is treated as a tunable parameter. In our experiments we set it to $\gamma = \frac{3}{4}$, which is proposed by Lindeberg in [16], and we compute the ridge strengths from $t = 1$ to $t = 30$.

In the scale space approach, ridge lines in 2D images sweep out surfaces in three-dimensional scale space, and curves on these surfaces are found along which ridge

strength is locally maximal with respect to diffusion time t. This means that, in general, different points on a ridge will correspond to different diffusion times. Thus, the algorithm that reconstructs the ridge geometry needs to operate on the full scale space, which can be challenging to implement especially when dealing with the four-dimensional scale space resulting from three-dimensional input images [15].

In contrast to this, the idea behind our filter is to run it as a pre-process that creates, at some fixed diffusion time, an image in which creases at a fine scale are preserved, while those at higher scales have been sufficiently enhanced so that their geometry can be reliably extracted even by an algorithm that operates on a single image, and does not have specific mechanisms for dealing with features at multiple scales [19].

In a similar spirit, prior work by Barakat et al. [2] approximated Lindeberg's approach by pre-computing derivatives at a locally optimal scale, and using them as input to a crease detection algorithm that would otherwise be able to reliably extract features only at a single fixed scale. The benefit of such an approach is that it greatly simplifies, and reduces the time and memory requirements of, the final crease detection. The main limitation is that it only considers a single optimal scale at each location while, in rare cases, a full scale space approach might indicate spatially intersecting creases at different scales.

4 Experimental Results

We compare our multi-scale anisotropic fourth-order diffusion (MAFOD) to crease enhancement diffusion (CED) [21], vesselness enhancement diffusion (VED) [4], isotropic fourth-order diffusion (IFOD) [17], the anisotropic fourth-order diffusion by Hajiaboli [9], bilateral, and a multi-scale Gaussian filter.

The multi-scale Gaussian filter is defined to approximate Lindeberg's scale selection, as described in Section 3.6. It first selects an optimal scale for each pixel, by finding the diffusion times t at which $R(u_\sigma)$ from Equation (19) is maximal. Then, the intensity of each pixel in the output image is obtained by convolving the input image with a Gaussian at the locally optimal scale $\sigma = \sqrt{2t}$ that is then normalized between $[0, 1]$. The normalization is necessary to compensate for the intesity range shrinkage after Gaussian blurring.

Localization accuracy is quantified by using a marching squares approach analogous to [19] for extracting crease lines in the filtered images, and measuring their distance to a ground truth. For each crease line segment in the ground truth, a corresponding segment in the reconstruction is selected by picking the one with minimum Hausdorff distance [12] in a neighborhood around the ground truth line segment. This neighborhood is set to six pixels for the experiments on synthetic data, and to ten pixels for the experiments on real data. The average Euclidean distance $\mathcal{E}$ between the ground truth and the corresponding reconstruction is then used to quantify the accuracy of vessel locations in the filtered image.

4.1 Confirming Theoretical Properties

Our filter has been designed to improve localization accuracy while accounting for creases at multiple scales and being rotationally invariant. Results on a simple simulated image with three concentric ridges of different sizes, which is contaminated with zero-mean Gaussian noise with a signal to noise ratio SNR = 6.81, verify that these design goals are met.

In Figure 4, our MAFOD filter restores ridge locations most accurately as assessed both by visual inspection and Euclidean distance $\mathcal{E}$. MAFOD outperforms CED, IFOD, Hajiaboli and bilateral filtering since it accounts for different scales. On the other hand, the curvature enhancement of our filter, which is not part of multiscale VED or Gaussian filters, clearly makes it easier for the ridge extraction algorithm to localize the centerline, especially in the largest ridge. IFOD does perform curvature enhancement but, due to its isotropic nature, it is not effectively guided to act specifically across the ridge. As it is obvious on the largest circle, the multi-scale Gaussian filter leads to ridge displacement. The result of the anisotropic fourth-order filter by Hajiaboli clearly illustrates the fact that it was designed to preserve edges, not to enhance creases.

For a fair comparison, image evolution of all filters, except for multi-scale Gaussian and bilateral filters, was stopped when the ℓ_2 difference between the filtered image and the noise-free ground truth was minimized. For the MAFOD filter, scales σ and vesselness threshold θ are the same as for VED, $\sigma = \{0.5, 1.0, ..., 8.5, 9.0\}$, $\theta = 0.2$. Other parameters are $\lambda = 0.005$ for MAFOD, IFOD, and $\sigma = 1.0$ for IFOD; for Hajiaboli, $\lambda = 0.01$; for CED, $\sigma = 2.0$ and it is set to enhance both ridges and valleys; for single-scale Gaussian smoothing, $\sigma = 1.25$. For the MAFOD filter, $\tau = 0.05$. For other fourth-order equations $\tau = 0.03$ and for the second-order diffusion equations such as the VED and CED filters, $\tau = 0.2$; for the bilateral filter $\sigma_{spatial} = 3.0$ and $\sigma_{range} = 1.0$.

4.2 Simulated Vessel Occlusion

Figure 5 shows a second image, simulating an occluded vessel, and corrupted with Gaussian noise with SNR = 6.40. Our MAFOD filter leads to the most accurate localization in terms of Euclidean error $\mathcal{E}$. In particular, we observed that VED widens the occlusions. They are better preserved by our filter, which we set to enhance both ridges and valleys on all simulated images.

Again, an amount of smoothing that minimized ℓ_2 error was used for all filters except for multi-scale Gaussian and bilateral filters. The parameters for VED and MAFOD are $\sigma = \{0.5, 1.0, 1.5, 2.0, 2.5, 3.0\}$, and $\theta = 0.35$; $\lambda = 0.017$ for MAFOD; for CED, $\sigma = 1.0$; for the bilateral filter, $\sigma_{\text{spatial}} = 1.5$ and $\sigma_{\text{range}} = 1.0$. For the numerical solver, we set $\tau = 0.05$ for the fourth-order equation and for second-order equations, $\tau = 0.2$.

4.3 Real Vessel Tree

To demonstrate our filter on a real-world example, we applied it to several ROIs from a fundus image, on which one of our co-authors (MWMW), who is an ophthalmologist and was blinded to our results, manually marked the exact vessel locations to provide

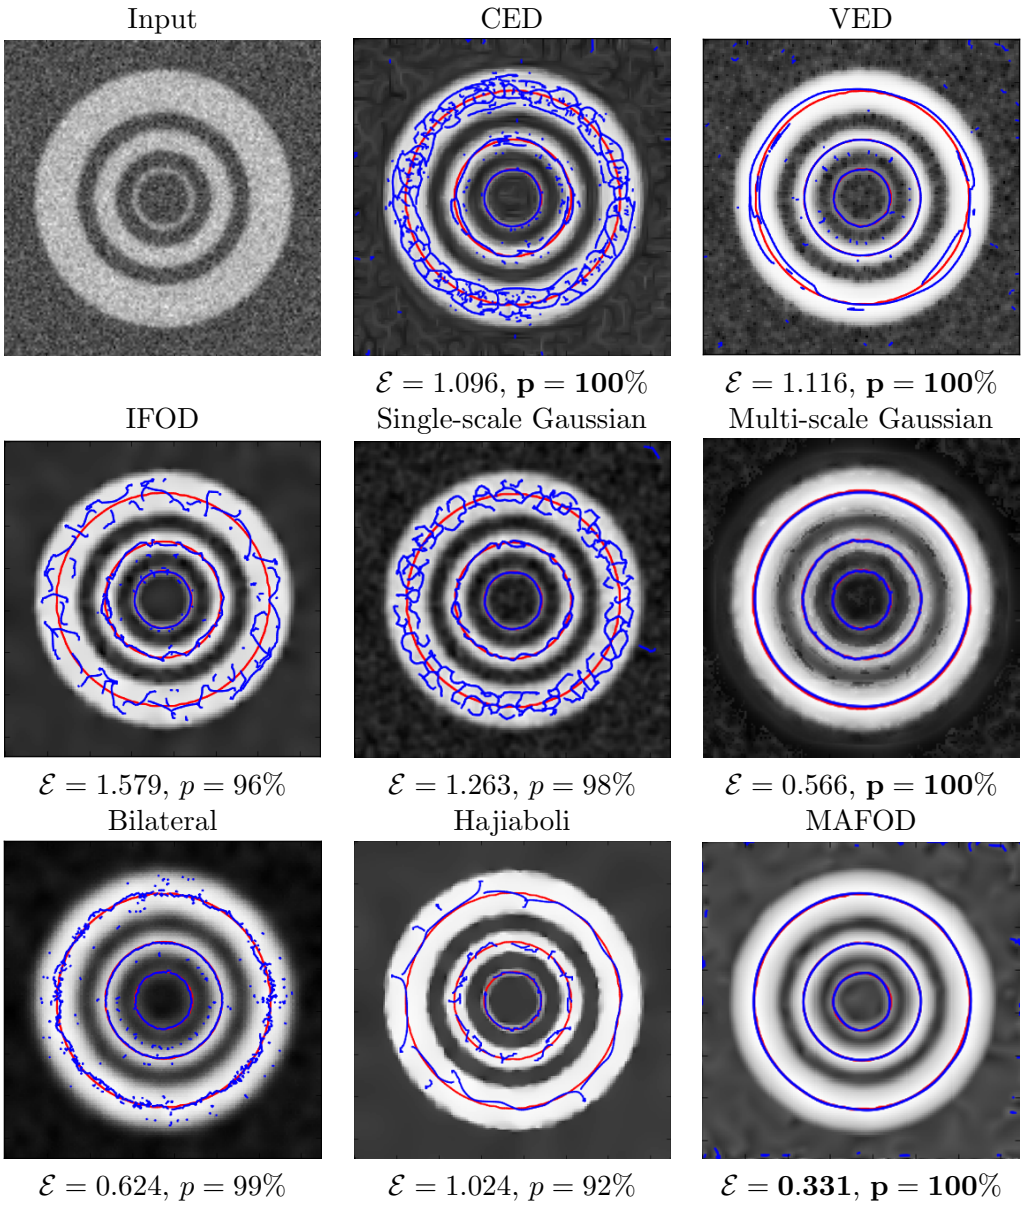

Figure 4: Red curves show the ground truth vessel location, while blue curves show the location reconstructed from the filtered noisy image. Our MAFOD filter restores ridge locations from the noisy image with ridges of different scales better than other filters. Value p shows the percentage of ground truth for which a corresponding ridge was detected from the filtered images while computing $\mathcal{E}$.

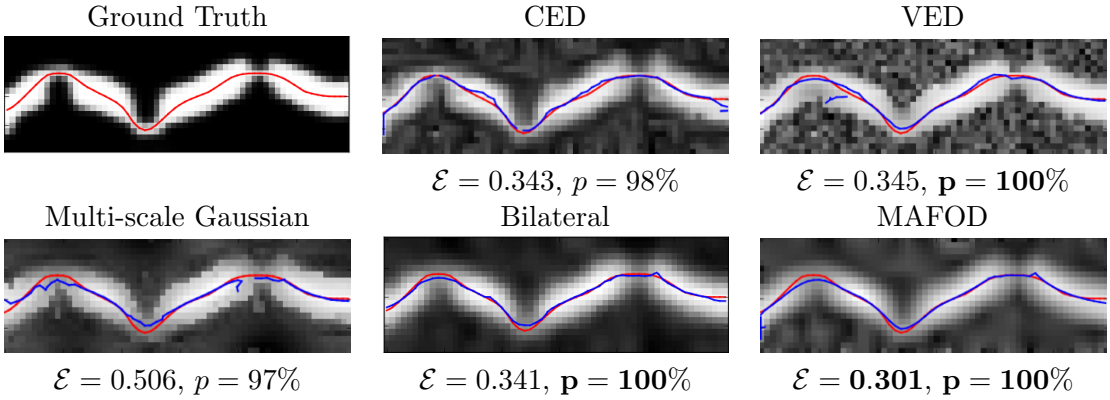

Figure 5: In a simulated occluded vessel, restored (blue) curves again best match the (red) ground truth locations in case of our MAFOD filter. In addition, MAFOD better preserves the occlusions than VED.

Filter	VED	Bilateral	Multi-scale Gaussian	MAFOD
Time (*sec*)	1251	0.16	2.24	6.51

Table 1: The average filtering time for a single ROI of size 200×200 pixels in Figure 6.

a ground truth for comparison. Results in Figure 6 show that our MAFOD filter outperforms VED, multi-scale Gaussian and bilateral filters in restoring vessel locations. In ROI 3, two vessels run close to each other, and are erroneously connected in the VED filtered image, even though not in the reconstructed curves. Our MAFOD filter successfully avoids this.

Even though vessels generally appear dark (i.e., as valleys) in these images, the larger ones exhibit a thin ridge at their center. This leads to an incorrect double response in single-scale filters as shown for the bilateral filter. CED and IFOD filters suffer from similar problems (results not shown).

For each filter separately, we carefully tuned the parameters for optimum results. For the MAFOD filter, we set $\sigma = \{0.2, 0.3, 0.5, 1.0, ..., 6.5, 7.0\}$, $\lambda = 0.005$, $\theta = 0.13$, and used a FED scheme with stopping time 12, cycle number 2 and $\tau_{\max} = 0.05$. For the VED filter, an explicit Euler scheme is used with 600 iterations and $\tau = 0.2$, and the same parameters for scale selection as for MAFOD; for the bilateral filter, $\sigma_{\text{range}} = 0.3$ and $\sigma_{\text{spatial}} = 3.0$; for the multi-scale Gaussian filter an additional Gaussian smoothing with kernel size $\sigma = 2$ is applied to the filtered image to blur out discontinuities from scale selection and thus achieve an even better result. The computational effort of all filters is reported in Table 1.

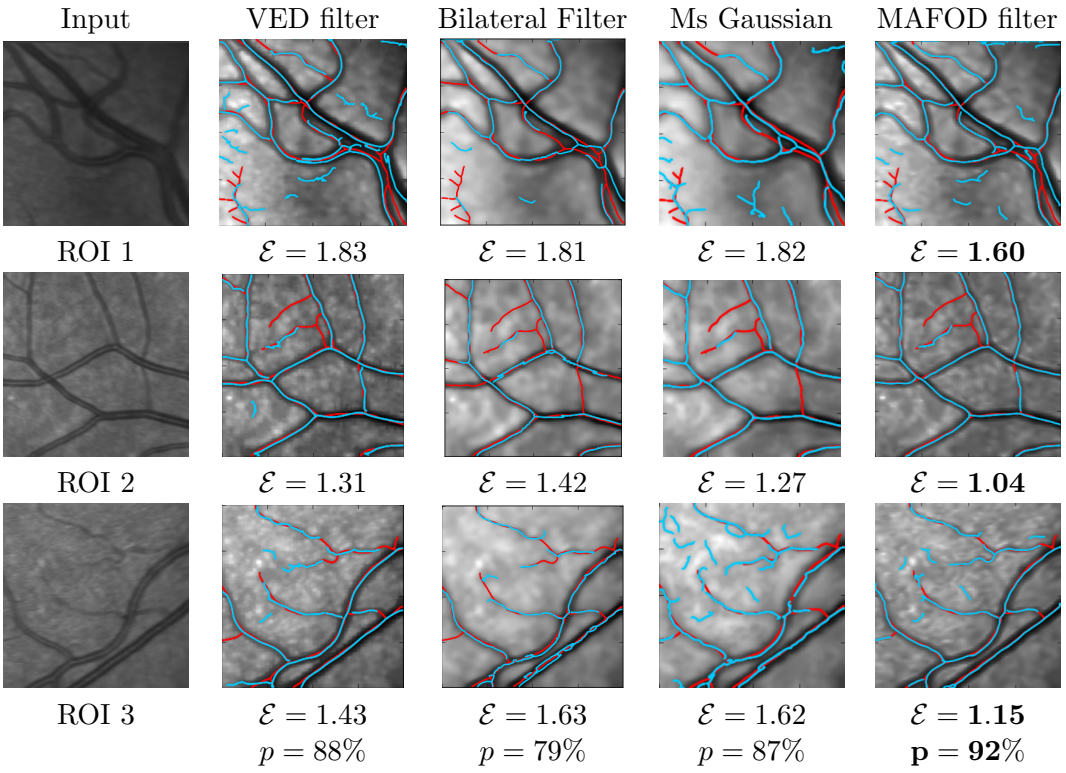

Figure 6: In three ROIs of a fundus image, reconstructed vessel locations (blue) best match a manually marked ground truth (red) when our MAFOD filter is used.

5 Conclusion

We have proposed a new multi-scale fourth order anisotropic diffusion (MAFOD) filter to enhance ridges and valleys in images. It uses a fourth order diffusion tensor which smoothes along creases, but sharpens them in the perpendicular direction, and optionally enables enhancing either ridges or valleys only. Our results indicate that the curvature enhancing properties of fourth-order diffusion allow our filter to better restore the exact crease locations than traditional methods. In addition, we found that our filter better preserves vessel occlusion.

In the future, we would like to extend our 2-D filter to 3-D images, and to better handle crossings and bifurcations [10].

References

[1] R. Annunziata, A. Kheirkhah, P. Hamrah, and E. Trucco. Scale and curvature invariant ridge detector for tortuous and fragmented structures. In *Proc. Medical*

Image Computing and Computer-Assisted Intervention (MICCAI), Part III, volume 9351 of *LNCS*, pages 588–595. Springer, 2015.

[2] S. Barakat, N. Andrysco, and X. Tricoche. Fast extraction of high-quality crease surfaces for visual analysis. *Computer Graphics Forum*, 30(3):961–970, 2011.

[3] P. J. Basser and S. Pajevic. Spectral decomposition of a 4th-order covariance tensor: Applications to diffusion tensor MRI. *Signal Processing*, 87:220–236, 2007.

[4] C. Cañero and P. Radeva. Vesselness enhancement diffusion. *Pattern Recognition Letters*, 24(16):3141–3151, 2003.

[5] M. Descoteaux, D. L. Collins, and K. Siddiqi. A geometric flow for segmenting vasculature in proton-density weighted mri. *Medical image analysis*, 12(4):497–513, 2008.

[6] S. Didas, J. Weickert, and B. Burgeth. Properties of higher order nonlinear diffusion filtering. *Journal of Mathematical Imaging and Vision*, 35(3):208–226, 2009.

[7] D. H. Eberly and S. M. Pizer. Ridge flow models for image segmentation. In *Medical Imaging 1994*, pages 54–64. International Society for Optics and Photonics, 1994.

[8] A. F. Frangi, W. J. Niessen, K. L. Vincken, and M. A. Viergever. Multiscale vessel enhancement filtering. In *Proc. Medical Image Computing and Computer-Assisted Interventation (MICCAI)*, volume 1496 of *LNCS*, pages 130–137. Springer, 1998.

[9] M. R. Hajiaboli. An anisotropic fourth-order diffusion filter for image noise removal. *International Journal of Computer Vision*, 92(2):177–191, 2011.

[10] J. Hannink, R. Duits, and E. Bekkers. Crossing-preserving multi-scale vesselness. In *Proc. Medical Image Computing and Computer-Assisted Intervention (MICCAI), Part II*, volume 8674 of *LNCS*, pages 603–610. Springer, 2014.

[11] R. M. Haralick, L. T. Watson, and T. J. Laffey. The topographic primal sketch. *Int'l Journal of Robotics Research*, 2(1):50–72, 1983.

[12] D. P. Huttenlocher, G. Klanderman, and W. J. Rucklidge. Comparing images using the Hausdorff distance. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 15(9):850–863, 1993.

[13] T. Jerman, F. Pernuš, B. Likar, and Ž. Špiclin. Enhancement of vascular structures in 3D and 2D angiographic images. *IEEE Trans. on Medical Imaging*, 35(9):2107–2118, 2016.

[14] G. Kindlmann, D. Ennis, R. Whitaker, and C.-F. Westin. Diffusion tensor analysis with invariant gradients and rotation tangents. *IEEE Trans. on Medical Imaging*, 26(11):1483–1499, 2007.

[15] G. Kindlmann, R. San José Estépar, S. M. Smith, and C.-F. Westin. Sampling and visualizing creases with scale-space particles. *IEEE Trans. on Visualization and Computer Graphics*, 15(6):1415–1424, 2009.

[16] T. Lindeberg. Edge detection and ridge detection with automatic scale selection. *International Journal of Computer Vision*, 30(2):117–156, 1998.

[17] M. Lysaker, A. Lundervold, and X.-C. Tai. Noise removal using fourth-order partial differential equation with applications to medical magnetic resonance images in space and time. *IEEE Trans. on Image Processing*, 12(12):1579–1590, 2003.

[18] P. Perona and J. Malik. Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 12(7):629–639, 1990.

[19] T. Schultz, H. Theisel, and H.-P. Seidel. Crease surfaces: From theory to extraction and application to diffusion tensor MRI. *IEEE Trans. on Visualization and Computer Graphics*, 16(1):109–119, 2010.

[20] T. Schultz, J. Weickert, and H.-P. Seidel. A higher-order structure tensor. In D. H. Laidlaw and J. Weickert, editors, *Visualization and Processing of Tensor Fields – Advances and Perspectives*, pages 263–280. Springer, 2009.

[21] A. F. Sóle, A. López, and G. Sapiro. Crease enhancement diffusion. *Computer Vision and Image Understanding*, 84(2):241–248, 2001.

[22] J. Weickert. *Anisotropic diffusion in image processing.* Teubner Stuttgart, 1998.

[23] J. Weickert, S. Grewenig, C. Schroers, and A. Bruhn. Cyclic schemes for PDE-based image analysis. *International Journal of Computer Vision*, pages 1–25, 2016.